# Dynamic Controllability of Conditional Simple Temporal Networks is PSPACE-complete

Massimo Cairo
*Mathematics Department*
*University of Trento*
*Trento, Italy*
*massimo.cairo@unitn.it*

Romeo Rizzi
*Computer Science Department*
*University of Verona*
*Verona, Italy*
*romeo.rizzi@univr.it*

arXiv:1608.08545v1 [cs.DS] 30 Aug 2016

***Abstract*—Even after the proposal of various solution algorithms, the precise computational complexity of checking whether a Conditional Temporal Network is Dynamically Controllable had still remained widely open. This issue gets settled in this paper which provides constructions, algorithms, and bridging lemmas and arguments to formally prove that: (1) the problem is PSPACE-hard, and (2) the problem lies in PSPACE.**

## I. Introduction

In temporal planning and scheduling, a *Simple Temporal Network* (STN) [1] consists of a set of tasks to be scheduled on the time line, and a set of constraints of the form $Y - X \leq \delta$, with $\delta \in \mathbb{R}$, i.e., limiting the difference between the execution times of tasks $X$ and $Y$. The STN is said to be consistent if it admits a schedule of its tasks that satisfies all the constraints. Some variants of the STN model have been proposed in the literature to allow and represent some forms of *contingency*, that is, the presence of parameters which are unknown to the planner. For example, a *Conditional Simple Temporal Network* (CSTN) [2] comprises also a set of unknown propositional variables, and some of the tasks and constraints in the network are to be taken into account only for specific values of these variables.

When contingency is present, such as in CSTNs, the notion of consistency is replaced by the notion of *controllability*, which comes in three flavors: weak, strong, and dynamic [3], [2]. In all the three variants the question is whether the planner is able to provide a schedule that satisfies the constraints; the difference is in how and when the value of the unknown parameters is disclosed to the planner. In the weak controllability, the parameters are revealed before the execution of the plan, so the schedule can be decided once the value of all the variables (scenario) has been specified, and the main question is deciding whether a feasible scheduling exists for all possible scenarios. In the strong controllability, these values are revealed only after the execution of the schedule, so we need a single schedule that works for every scenario. In the dynamic controllability, the unknown parameters are revealed progressively *during* the execution of the plan, as a consequence of actions performed by the planner. In the case of CSTNs, each propositional variable is associated with an observation task, and its value is revealed precisely when the corresponding observation task is executed. Here we look for a *dynamic execution strategy*: a schedule of the tasks that gets dynamically decided depending on the partial scenario progressively observed, such that, whatever scenario possibly emerges, all the constraints pertinent to that scenario will be respected. The dynamic controllability[1] decision problem for CSTNs (CSTN-DC) asks to check whether a given CSTN is dynamically controllable, and is a major algorithmic problem associated to CSTNs.

It is known that the consistency of STNs can be decided in polynomial time, by interpreting the network as a weighted graph and applying the Floyd-Warshall All-Pairs Shortest Path algorithm [1]. However, the presence of contingency might change drastically the algorithmic nature of the problem. Especially for the dynamic controllability, which introduces an alternation of quantifiers $\exists$ (for the choices of the planner) and $\forall$ (for the revealed parameters), an increase in complexity is expected [3]. In the case of *Simple Temporal Networks with Uncertainty* (STNUs), another variant of STNs with contingency (with which a first controllability issue was posed), the dynamic controllability had been conjectured to be PSPACE-complete [3]. Subsequently, it was proven to actually lie in P [4].

For CSTNs, determining the right complexity class of dynamic controllability is still a widely open question. Deciding their weak controllability has been proven to be coNP-complete [2], [5] and, since weak controllability can easily be reduced to a special case of dynamic controllability, then CSTN-DC is at least coNP-hard [5]. A

This work was supported by the Department of Computer Science, University of Verona, under PhD grant "Computational Mathematics and Biology".

[1]The term "dynamic consistency" is sometimes used in the context of CSTNs. We prefer to use "dynamic controllability" to emphasize the active role of the planner and to match the name used in the literature for other type of temporal networks, such as STNUs.

first complete algorithmic solution had been proposed in [6] by reducing CSTN-DC to a time automaton game of high complexity. Later, a complete constraint propagation algorithm was achieved with much better performances in practice [7], based on the sound constraint-propagation rules provided in [8], [9], [10] for the more general setting of Conditional Simple Temporal Networks with Uncertainty (CSTNU). In [5], a worst-case upper bound is obtained, thanks to an algorithm which requires singly-exponential time and memory. To the best of our knowledge, no better bounds are known in the literature.

In this work, we settle this question by sharply improving both the lower and the upper bound. After providing the background notions and first basic facts as common to both developments in Section II, a reduction from Quantified 3-SAT to CSTN-DC is proposed in Section III, which proves the latter to be PSPACE-hard. Then, in Section IV, the first algorithm that solves CSTN-DC using only polynomial memory is given, hence showing that CSTN-DC lies in PSPACE. Sections III and IV can be read independently. Taken together, their negative and positive results show that CSTN-DC is PSPACE-complete, i.e., the natural complexity class for the dynamic controllability issue for CSTNs is PSPACE.

## II. Background

In this section, Conditional Simple Temporal Networks (CSTNs) and their dynamic controllability are formally defined. The definitions are taken from [7].

### A. Simple Temporal Networks (STNs)

**Definition 1** (Temporal variables, tasks, constraints)**.** Let $\mathcal{T}$ be a finite set of real-valued *temporal variables*. Each variable $X \in \mathcal{T}$ represents the execution time of a task, also denoted with $X$. In the following, we use the terms temporal variable and task interchangeably. A *binary difference constraint* over $\mathcal{T}$ is a constraint of the form $Y - X \leq \delta$, for $X, Y \in \mathcal{T}$ and $\delta \in \mathbb{R}$. In this paper, *constraint* always denotes a binary difference constraint. The constraint $Y - X \leq \delta$ can also be expressed, equivalently, as $X - Y \geq -\delta$, $Y \leq X + \delta$ or $X \geq Y - \delta$, as it is more convenient in the context.

**Definition 2** (Schedule, satisfied constraints)**.** A *schedule* over $\mathcal{T}$ is a total assignment $\psi \colon \mathcal{T} \to \mathbb{R}$ of the temporal variables in $\mathcal{T}$. We write $[\psi]_X$ instead of $\psi(X)$ to denote the value assigned by the schedule $\psi$ to the variable $X \in \mathcal{T}$. A schedule *satisfies* a constraint $Y - X \leq \delta$ if $[\psi]_Y - [\psi]_X \leq \delta$.

**Definition 3** (Simple Temporal Network)**.** A *Simple Temporal Network* (STN) is a pair $(\mathcal{T}, \mathcal{C})$ where $\mathcal{T}$ is a set of temporal variables and $\mathcal{C}$ is a set of constraints over $\mathcal{T}$.

**Definition 4** (Feasible schedule)**.** A schedule $\psi$ over $\mathcal{T}$ is *feasible* for $(\mathcal{T}, \mathcal{C})$ if $\psi$ satisfies all the constraints in $\mathcal{C}$.

### B. Conditional Simple Temporal Networks (CSTNs)

**Definition 5** (Propositional variables, labels)**.** Let $\mathcal{P}$ be a set of propositional (boolean) variables. A *label* $\ell$ over $\mathcal{P}$ is a boolean formula $\ell = l_1 \wedge \cdots \wedge l_k$, obtained as conjunction of positive or negative literals $l_i \in \{p_i, \neg p_i\}$ on distinct variables $p_i \in \mathcal{P}$. The empty label is denoted with $\lambda$ and always evaluates to true. Let $\mathcal{P}^*$ denote the set of labels over $\mathcal{P}$ (including $\lambda$).

**Definition 6** (Scenario, label evaluation)**.** A *scenario* $s$ over $\mathcal{P}$ is a total assignment of the propositional variables $s \colon \mathcal{P} \to \{0, 1\}$ where 0 means *false* and 1 means *true*. Let $\Sigma_{\mathcal{P}}$ denote the set of all the scenarios over $\mathcal{P}$. We write $s \vDash \ell$ if the label $\ell$ evaluates to true under the interpretation given by $s$.

**Definition 7** (Conditional Simple Temporal Network)**.** A *Conditional Simple Temporal Network* (CSTN) is a tuple $\Gamma = (\mathcal{T}, \mathcal{P}, \mathcal{C}, \mathcal{L}, \mathcal{OT}, \mathcal{O})$ where

- $\mathcal{T}$ is a finite set of *temporal variables* or *tasks*,
- $\mathcal{P}$ is a finite set of *propositional variables*,
- $\mathcal{C}$ is a finite set of *labeled constraints* $(Y - X \leq \delta, \ell)$, where $Y - X \leq \delta$ is a constraint over $\mathcal{T}$ and $\ell \in \mathcal{P}^*$ is a label,
- $\mathcal{L} \colon \mathcal{T} \to \mathcal{P}^*$ if a function that assigns a label $\mathcal{L}(X)$ to each task $X \in \mathcal{T}$,
- $\mathcal{OT} \subseteq \mathcal{T}$ is the set of *observation tasks*,
- $\mathcal{O} \colon \mathcal{P} \to \mathcal{OT}$ is a bijection that associates each propositional variable $p \in \mathcal{P}$ to a unique observation task $\mathcal{O}(p)$.

A task $X \in \mathcal{T}$ has to be executed only in those scenarios $s \in \Sigma_{\mathcal{P}}$ such that $s \vDash \mathcal{L}(X)$, and each constraint $(Y - X \leq \delta, \ell) \in \mathcal{C}$ has to be satisfied if $s \vDash \ell$. Since the constraint $Y - X \leq \delta$ only makes sense if both $X$ and $Y$ get executed, we require the following well-definedness property.

**Definition 8** (Restriction WD1)**.** A CSTN satisfies the restriction *WD1* if, for every labeled constraint $(Y - X \leq \delta, \ell) \in \mathcal{C}$, we have $\ell \Rightarrow \mathcal{L}(X) \wedge \mathcal{L}(Y)$.

In the following, *WD1* is incorporated in the definition of CSTN, i.e., it is assumed that any CSTN satisfies this restriction.

*Remark* 1. Tsamardinos et al. [2] discussed some supplementary *reasonability assumptions* that any well-defined CSTN must satisfy. Subsequently, those conditions have been analyzed and formalized in [8] introducing the three restrictions *WD1*, *WD2*, and *WD3*. The restriction *WD1* has already been discussed. The restrictions *WD2* and *WD3* relate the labels on tasks and constraints with the labels on observation tasks. We avoid entering into the fine details regarding them, and we rather provide both of our results in their strongest and most general form, as follows. First, the reduction in our PSPACE-hardness proof constructs only CSTNs that comply with all three restrictions vacuously,

having no labels on the tasks. Second, neither *WD2* nor *WD3* are required as preconditions for the applicability of our PSPACE algorithm.[2]

### C. Dynamic controllability of CSTNs

**Definition 9** (Projection)**.** The *projection* of a CSTN over a scenario $s$ is the STN $\Gamma_s = (\mathcal{T}_s, \mathcal{C}_s)$ where:

- $\mathcal{T}_s = \{X \in \mathcal{T} \mid s \vDash \mathcal{L}(X)\}$
- $\mathcal{C}_s = \{Y - X \leq \delta \mid (Y - X \leq \delta, \ell) \in \mathcal{C} \text{ and } s \vDash \ell\}$.

**Definition 10** (Execution strategy, viable)**.** Let $\Psi_{\mathcal{T}}$ denote the set of schedules $\psi$ over any subset $\mathcal{T}' \subseteq \mathcal{T}$. For a schedule $\psi \in \Psi_{\mathcal{T}}$ over $\mathcal{T}' \subseteq \mathcal{T}$, let $\mathrm{Dom}(\psi) = \mathcal{T}'$. An *execution strategy* for $\Gamma$ is a function $\sigma\colon \Sigma_{\mathcal{P}} \to \Psi_{\mathcal{T}}$ that maps each scenario $s \in \Sigma_{\mathcal{P}}$ to a schedule $\sigma(s)$ for $\Gamma_s$ (i.e. $\mathrm{Dom}(\sigma(s)) = \mathcal{T}_s$). An execution strategy $\sigma$ is *viable* if, for every scenario $s \in \Sigma_{\mathcal{P}}$, the schedule $\sigma(s)$ is feasible for $\Gamma_s$.

**Definition 11** (Partial scenario, history)**.** A *partial scenario* over $\mathcal{P}$ is a partial assignment $h\colon \mathrm{Dom}(h) \to \{0,1\}$ of the propositional variables, where $\mathrm{Dom}(h) \subseteq \mathcal{P}$. Given $\sigma$, a scenario $s$, and a time point $t \in \mathbb{R}$, the *history* at $t$ in the scenario $s$ with the strategy $\sigma$ is the partial scenario $\mathrm{Hist}(t, s, \sigma)$ where $\mathrm{Dom}(\mathrm{Hist}(t, s, \sigma)) = \{p \in \mathcal{P} \mid [\sigma(s)]_{\mathcal{O}(p)} < t\}$ and $\mathrm{Hist}(t, s, \sigma)(p) = s(p)$ for every $p \in \mathrm{Dom}(\mathrm{Hist}(t, s, \sigma))$.

**Definition 12** (Dynamic execution strategy)**.** An execution strategy $\sigma$ is *dynamic* if, for any scenarios $s, s' \in \Sigma_{\mathcal{P}}$ and time variable $X \in \mathcal{T}_s$, letting $t = [\sigma(s)]_X$, if $\mathrm{Hist}(t, s, \sigma) = \mathrm{Hist}(t, s', \sigma)$, then $X \in \mathcal{T}_{s'}$ and $[\sigma(s')]_X = t$.

**Definition 13** (Dynamic controllability)**.** A CSTN is *dynamically controllable* if it admits a dynamic viable execution strategy. The *dynamic controllability* decision problem (CSTN-DC) asks to check whether a given CSTN $\Gamma$ is dynamically controllable or not.

The following definition and lemma state a useful characterization of dynamic execution strategies: if two scenarios differ in only one propositional variable, then a dynamic execution strategy behaves in the same way in the two scenarios until that propositional variable is observed. This property has been also proven in [7, Theorem 1], and will be used several times in this paper to exploit the fact that an execution strategy is dynamic.

**Definition 14.** Given a scenario $s \in \Sigma_{\mathcal{P}}$, a propositional variable $p \in \mathcal{P}$ and $v \in \{0,1\}$, let $s[v/p]$ be the scenario obtained from $s$ by changing the value of the variable $p$ to $v$, i.e., $s[v/p](p) = v$ and $s[v/p](q) = s(q)$ for $q \in \mathcal{P} \setminus \{p\}$.

**Lemma 1.** *Let $\sigma$ be a dynamic execution strategy. Let $s \in \Sigma_{\mathcal{P}}$, $p \in \mathcal{P}$, $v \in \{0,1\}$, and consider the scenario $s' = s[v/p]$. For any $t \leq [\sigma(s)]_{\mathcal{O}(p)}$, the following properties hold:*

*(a)* $\mathrm{Hist}(t, s, \sigma) = \mathrm{Hist}(t, s', \sigma)$,
*(b)* $[\sigma(s)]_X \leq t \iff [\sigma(s')]_X \leq t$ *for every* $X \in \mathcal{T}$,
*(c)* $[\sigma(s)]_X = t \iff [\sigma(s')]_X = t$ *for every* $X \in \mathcal{T}$

*and in particular* $[\sigma(s)]_{\mathcal{O}(p)} = [\sigma(s')]_{\mathcal{O}(p)}$.

*Proof:* Observe that we only need to check the properties for those values of $t$ that appear somewhere in $\sigma$. Hence, we can proceed by induction. The properties clearly hold for a sufficiently small $t$ (less than any value in $\sigma$). Take $t < t' \leq [\sigma(s)]_{\mathcal{O}(p)}$, such that $[\sigma(s)]_X, [\sigma(s')]_X \notin (t, t')$ for every $X \in \mathcal{T}$, and assume that the properties hold for $t$. Observe that, for $s'' \in \{s, s'\}$, $\mathrm{Dom}(\mathrm{Hist}(t', s'', \sigma)) = \mathrm{Dom}(\mathrm{Hist}(t, s'', \sigma)) \cup \{q \in \mathcal{P} \mid [\sigma(s'')]_{\mathcal{O}(q)} = t\}$ and, since $[\sigma(s)]_{\mathcal{O}(q)} = t \iff [\sigma(s')]_{\mathcal{O}(q)} = t$ for $q \in \mathcal{P}$, then $\mathrm{Dom}(\mathrm{Hist}(t', s, \sigma)) = \mathrm{Dom}(\mathrm{Hist}(t', s', \sigma))$. Moreover, $p \notin \mathrm{Dom}(\mathrm{Hist}(t', s, \sigma))$ since $t < [\sigma(s)]_{\mathcal{O}(p)}$ by assumption. Hence, $\mathrm{Hist}(t', s, \sigma) = \mathrm{Hist}(t', s', \sigma)$. The other two properties for $t'$ are a direct consequence of the fact that $\sigma$ is dynamic, and this concludes the induction. ■

## III. PSPACE-HARDNESS

In this section, we prove that CSTN-DC is PSPACE-hard by showing a reduction from Quantified 3-SAT (Q3SAT).

We are given a Q3SAT formula $\Phi = \exists x_1 \forall y_1 \cdots \exists x_n \forall y_n \varphi$ where $\varphi$ is a 3CNF over the propositional variables $x_1, y_1, \ldots, x_n, y_n$. I.e., $\varphi = \bigwedge_{j=1}^{m} (l_{j,1} \vee l_{j,2} \vee l_{j,3})$ and each literal $l_{j,k}$ is either a positive or a negated occurrence of one of the quantified variables. The formula $\Phi$ can be understood as a game in which the existential player and the universal player decide in turn the value of the variables $x_1, y_1, x_2, y_2, \ldots, x_n, y_n$. The existential players wins if, when all the variables have been set, the formula $\varphi$ is satisfied by the chosen values. CSTNs can be also seen as games, where the planner plays against the nature, the first by scheduling the tasks, the second by choosing the value of the propositional variables as soon as they are observed. The planner wins if, eventually, the schedule he executes is feasible, and the CSTN is dynamically controllable if the planner has a winning strategy. This interpretation of both Q3SAT and CSTN-DC as two-player games underlies our proof of PSPACE-hardness.

We will describe a CSTN $\Gamma_\Phi$ which is dynamically controllable iff $\Phi \equiv$ true, that is, iff the existential player has a winning strategy for $\Phi$. It will be apparent that $O(\log(n+m))$ internal space suffices in order to construct $\Gamma_\Phi$ out from $\Phi$.

### A. Warm-up: the controller can choose some variables

Before addressing CSTN-DC, we consider a more general problem CSTN$^+$-DC. We define a CSTN$^+$ to be a CSTN in which the values of a subset $\mathcal{P}^+ \subseteq \mathcal{P}$ of the propositional variables are actually *decided* by the controller rather than

[2]The reader may observe that, without *WD2* and *WD3*, it is possible to have the corner case of a network which does not admit any dynamic execution strategy [8]. Such a network is considered not dynamically controllable since, in particular, it does not admit any *viable* dynamic execution strategy. No special handling of this case is needed.

by the nature, still each $p \in \mathcal{P}$ gets determined at the precise execution time of the corresponding disclosure task $\mathcal{O}(p)$. To ease our exposition, we first construct a CSTN$^+$ $\Gamma^+_\Phi$ which is dynamically controllable iff $\Phi \equiv$ true. This will be a much easier task, but helps in delivering the general idea of the reduction.

The CSTN$^+$ $\Gamma^+_\Phi$ contains all the variables $x_i, y_i$ as propositional variables, decided and observed respectively in tasks $X_i = \mathcal{O}(x_i)$ and $Y_i = \mathcal{O}(y_i)$. These tasks are subject to the unlabeled constraints $Y_i \geq X_i + 1$ $(i = 1, \ldots, n)$ and $X_{i+1} \geq Y_i + 1$ $(i = 1, \ldots, n-1)$. These constraints connect the tasks $X_1, Y_1, \ldots, X_n, Y_n$ in a chain which enforces that they are executed in the proper order. Then, we have two tasks $A$ and $B$ with the following constraints. For every $j = 1, \ldots, m$, we have a constraint $B \geq A + 1$ with label $\ell_j := \neg l_{j,1} \wedge \neg l_{j,2} \wedge \neg l_{j,3}$, defined as the negation of the $j$-th clause $(l_{j,1} \vee l_{j,2} \vee l_{j,3})$ of $\varphi$. Finally, there is an unlabeled constraint $A \geq B + 1$.

The network $\Gamma^+_\Phi$ is dynamically controllable iff $\Phi \equiv$ true. Indeed, if $\Phi \equiv$ true, the controller schedules $X_i$ and $Y_i$ at time $2i$ and $2i + 1$ respectively, and can choose the propositional value of $x_i$ depending on $y_1, \ldots, y_{i-1}$ in accordance to his winning strategy for $\Phi$. Finally, he schedules $B$ at time $2n + 2$ and $A$ at time $2n + 3$, and, since every clause of $\varphi$ is satisfied, none of the constraints $B \geq A + 1$ applies and all the other constraints are fulfilled. Conversely, assume $\Phi \equiv$ false. It is now nature that owns a winning strategy: since the controller is anyhow forced to reveal the variables in order, she can choose each propositional variable $y_i$ depending on $x_1, \ldots, x_i$ so that, for at least one $j \in \{1, \ldots, m\}$, the clause $(l_{j,1} \vee l_{j,2} \vee l_{j,3})$ evaluates to false. Hence, the constraints $A \geq B + 1$ and $B \geq A + 1$ necessarily lead to a conflict, no matter when the events $A$ and $B$ are scheduled.

### *B. Reduction for CSTNs*

The above toy reduction with $\Gamma^+_\Phi$ illustrates well the general framework, but relies on the strong assumption that the controller can *choose* the value of some of the propositional variables. In CSTNs, the controller cannot force the nature to choose a particular value for a propositional variable. However, he can put a lot of pressure on her to choose the value he wants. Indeed, we next describe a network (in the standard framework of CSTNs) that allows the controller to specify the value he desires for a variable $x_i$, by executing one among two particular actions, one for true and one for false. The network is built in such a way that, if the value actually chosen by the nature differs from the prescription of the planner, then he is able to schedule the rest of the network easily, satisfying all the remaining constraints. Thanks to this property, the nature is effectively obliged to choose the variables as specified by the planner, otherwise she is doomed to lose the match.

We begin with an informal description of our construction $\Gamma_\Phi$. Figure 1 shows an example of our construction for $n = 3$, and may help the reader in following the exposition. There are $n$ gadgets $G_1, \ldots, G_n$ connected in series. The $i$-th gadget $G_i$ involves the propositional variables $x_i, y_i$ (which are now normal variables chosen by the nature), and two extra variables $c_i^1$ and $c_i^0$. The purpose of $G_i$ is to let the controller choose the value of $x_i$, and then observe the value of $y_i$ chosen by the nature. The nodes of $G_i$ are $A_i$, $B_i$, $C_i^0 = \mathcal{O}(c_i^0)$, $C_i^1 = \mathcal{O}(c_i^1)$, $D_i$, $X_i = \mathcal{O}(x_i)$ and $Y_i = \mathcal{O}(y_i)$. Moreover, $G_i$ connects also to the nodes $A_{i+1}$ and $B_{i+1}$ which, for $i < n$, belong to the next gadget $G_{i+1}$, while $A_{n+1}$ and $B_{n+1}$ are two extra nodes at the end of the construction. It is here, between $A_{n+1}$ and $B_{n+1}$, that the $m$ clause constraints get lied down. For each $j = 1, \ldots, m$, we put a constraint $B_{n+1} - A_{n+1} \geq n + 1$ with label $\ell_j := \neg l_{j,1} \wedge \neg l_{j,2} \wedge \neg l_{j,3}$, defined as the negation of the $j$-th clause of $\varphi$, like in the toy reduction of the previous section.

Before describing the internals of each gadget, we show how they play together and we focus only on the tasks $A_i$ and $B_i$ for $i = 1, \ldots, n + 1$. Consider the constraint $B_i - A_i \geq i - 1$ for $1 \leq i \leq n + 1$, called "activation constraint". The gadget $G_i$ is "activated" if the $i$-th activation constraint is satisfied, i.e., if the task $B_i$ is executed at most $i - 1$ units of time after $A_i$. For the first gadget $G_1$, the activation constraint $B_1 - A_1 \leq 0$ is explicitly added to the network, without labels, enforcing the gadget $G_1$ to be always activated. Thanks to the internal structure of the gadgets, the activation constraint is then propagated from one gadget to the next, as long as the nature chooses the value of $x_i$ according to the prescription of the controller. If the nature always chooses $x_i$ according to the controller, then all the gadgets are activated and we end up with the propagated constraint $B_{n+1} - A_{n+1} \leq n$. At this point, the controller is able to schedule $A_{n+1}$ and $B_{n+1}$ if and only if all the clauses of $\varphi$ are satisfied, so that the constraints $B_{n+1} - A_{n+1} \geq n + 1$ labeled with $\ell_i$ are all void.

If instead the nature chooses for any $x_i$ the opposite value to the one prescribed, then the activation constraint on $A_i$ and $B_i$ is not propagated to $A_{i+1}$ and $B_{i+1}$, the following gadgets are not activated, and the controller is able to execute all the other tasks $B_{i'}$, $C_{i'}^h$ and $D_{i'}$ for $i' > i$ very far in the future, without violating any constraint.

We now describe the internal mechanism of each gadget. At the heart of $G_i$ there are the two constraints $D_i \leq B_i + 1$ and $D_i \geq A_i + (n + 2)$, labeled with $c_i^1 \wedge c_i^0$ and $\neg c_i^1 \wedge \neg c_i^0$ respectively. If the gadget is activated, then these two constraints cannot be satisfied together (since $B_i - A_i \leq i - 1 \leq n$). Hence, the controller has to observe either $c_i^1$ or $c_i^0$, in order to decide whether to execute $D_i$ early or not (see Lemma 3). Which of the two variables is observed specifies the desired value of $x_i$: so, if the planner wants $x_i$ to be true, then he should execute $C_i^1$, and if he wants $x_i$ to be

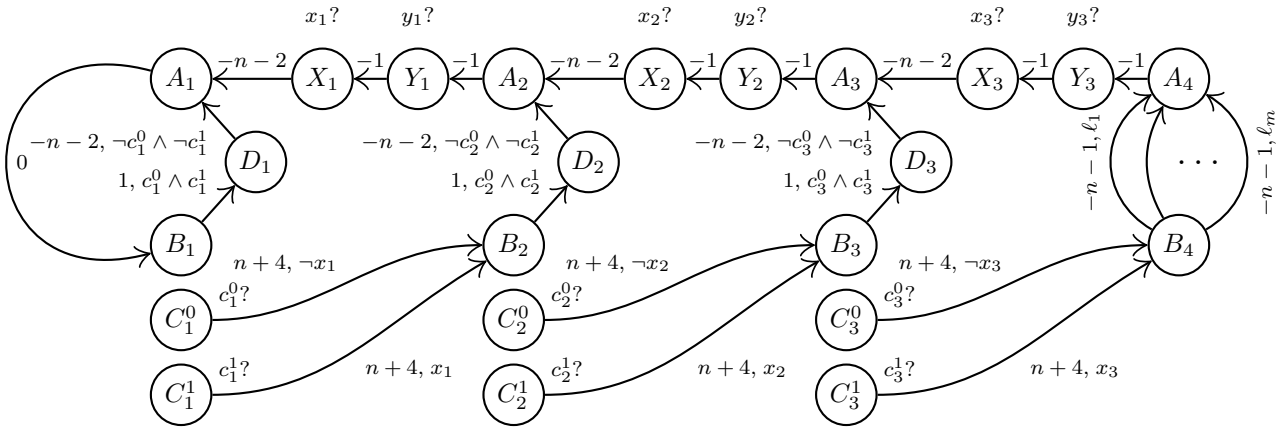
$x_1?$ $y_1?$ $x_2?$ $y_2?$ $x_3?$ $y_3?$
$A_1$ $X_1$ $Y_1$ $A_2$ $X_2$ $Y_2$ $A_3$ $X_3$ $Y_3$ $A_4$
$-n-2$ $-1$
$D_1$ $D_2$ $D_3$
$0$ $-n-2, \neg c_1^0 \wedge \neg c_1^1$ $1, c_1^0 \wedge c_1^1$
$-n-2, \neg c_2^0 \wedge \neg c_2^1$ $1, c_2^0 \wedge c_2^1$
$-n-2, \neg c_3^0 \wedge \neg c_3^1$ $1, c_3^0 \wedge c_3^1$
$-n-1, \ell_1$ $-n-1, \ell_m$ $\cdots$
$B_1$ $B_2$ $B_3$ $B_4$
$n+4, \neg x_1$ $n+4, \neg x_2$ $n+4, \neg x_3$
$c_1^0?$ $c_2^0?$ $c_3^0?$
$C_1^0$ $C_2^0$ $C_3^0$
$C_1^1$ $C_2^1$ $C_3^1$
$c_1^1?$ $c_2^1?$ $c_3^1?$
$n+4, x_1$ $n+4, x_2$ $n+4, x_3$

Figure 1. Reduction from Q3SAT to CSTN-DC for $n = 3$. A Q3SAT formula $\Phi = \exists x_1 \forall y_1 \exists x_2 \forall y_2 \exists x_3 \forall y_3 \bigwedge_{j=1}^{m} (l_{j,1} \vee l_{j,2} \vee l_{j,3})$ is transformed into the network $\Gamma_\Phi$ illustrated above, where nodes denote tasks and edges denote constraints. Specifically, a directed edge labeled $\delta, \ell$ from a node $N$ to a node $M$ denotes the labeled constraint $(M \leq N + \delta, \ell)$. The label $\ell_j$ is defined as the negation of the $j$-th clause of the 3CNF formula, i.e., $\ell_j \coloneqq \neg l_{j,1} \wedge \neg l_{j,2} \wedge \neg l_{j,3}$. The empty label $\lambda$ is omitted. A label $q?$ beside a node $Q$ indicates that $Q = \mathcal{O}(q)$.

false, he should execute $C_i^0$. A constraint $X_i \geq A_i + (n+2)$ ensures that the controller can observe $x_i$ only after choosing either $C_i^1$ or $C_i^0$. Then, the constraint $Y_i \geq X_i + 1$ allows the controller to observe $y_i$ only after $x_i$, and the constraint $A_{i+1} \geq Y_i + 1$ connects to the next gadget.

The propagation of the activation constraint $B_i - A_i \leq i - 1$ to the next gadget is achieved by the two constraints $B_{i+1} \leq C_i^1 + (n+4)$ and $B_{i+1} \leq C_i^0 + (n+4)$, labeled with $x_i$ and $\neg x_i$ respectively. In order for the propagation to take place, the nature has to choose $x_i$ to true if $C_i^1$ has been executed, and to false if $C_i^0$ has been executed (see Lemma 5).

The full construction is provided for reference in Figure 2, and it is illustrated in Figure 1 for $n = 3$.

**Lemma 2.** *If $\Phi \equiv$ true then $\Gamma_\Phi$ is dynamically controllable.*

*Proof:* Assume $\Phi \equiv$ true. This means that the existential player holds a winning strategy for $\Phi$. This strategy can be expressed as a function $f\colon \{0,1\}^* \to \{0,1\}$, such that $\varphi$ evaluates to *true* over all truth-assignments $s\colon \{x_1, y_1, \ldots, x_n, y_n\} \to \{0,1\}$ in which $s(x_i) = f(s(y_1), \ldots, s(y_{i-1}))$ for every $i = 1, \ldots, n$. Taking $f$ as reference, we provide a viable and dynamic execution strategy for $\Gamma_\Phi$.

Given a scenario $s\colon \mathcal{P} \to \{0,1\}$, let $h_i(s) = f(s(y_1), \ldots, s(y_{i-1}))$ for $i = 1, \ldots, n$. Then, define $b(s)$ to be the smallest index $i \in \{1, \ldots, n\}$ such that $s(x_i) \neq h_i(s)$, or $b(s) = n + 1$ if no such index $i$ exists. The value $b(s)$ represents the first index $i$ in which the nature does not follow the prescription of the controller in choosing the value of $x_i$, or $b(s) = n + 1$ if she copies until the end.

The execution strategy $\sigma$ is defined in Figure 3, where $b'(s) \coloneqq \min\{b(s), n\}$ and $\infty$ denotes a sufficiently large value, say, $\infty \coloneqq (n+4)(n+2)$.

$\mathcal{T} = \{A_i, B_i, C_i^0, C_i^1, D_i, X_i, Y_i\}_{i=1,\ldots,n} \cup \{A_{n+1}, B_{n+1}\}$
$\mathcal{P} = \{x_i, y_i, c_i^1, c_i^0\}_{i=1,\ldots,n}$,
$\mathcal{L}(N) = \lambda$ for every $N \in \mathcal{T}$,
$\mathcal{C}$ contains the following constraints:
  $(B_1 - A_1 \leq 0, \lambda)$,
  for $i = 1, \ldots, n$:
    $(D_i \leq B_i + 1, c_i^0 \wedge c_i^1)$,
    $(D_i \geq A_i + (n+2), \neg c_i^0 \wedge \neg c_i^1)$,
    $(X_i \geq A_i + (n+2), \lambda)$,
    $(Y_i \geq X_i + 1, \lambda)$,
    $(A_{i+1} \geq Y_i + 1, \lambda)$,
    $(B_{i+1} \leq C_i^0 + (n+4), \neg x_i)$,
    $(B_{i+1} \leq C_i^1 + (n+4), x_i)$,
  for $j = 1, \ldots, m$:
    $(B_{n+1} - A_{n+1} \geq n + 1, \neg l_{j,1} \wedge \neg l_{j,2} \wedge \neg l_{j,3})$,
$\mathcal{OT} = \{X_i, Y_i, C_i^0, C_i^1\}_{i=1,\ldots,n}$,
$\mathcal{O}(x_i) = X_i$, $\mathcal{O}(y_i) = Y_i$, $\mathcal{O}(c_i^0) = C_i^0$ and $\mathcal{O}(c_i^1) = C_i^1$, for $i = 1, \ldots, n$.

Figure 2. Construction of the CSTN $\Gamma_\Phi = (\mathcal{T}, \mathcal{P}, \mathcal{C}, \mathcal{L}, \mathcal{OT}, \mathcal{O})$ for a given Q3SAT formula $\Phi = \exists x_1 \forall y_1 \cdots \exists x_n \forall y_n \bigwedge_{j=1}^{m} (l_{j,1} \vee l_{j,2} \vee l_{j,3})$.

Notice that the value $[\sigma(s)]_N$ for a task $N \in \mathcal{T}$ and a scenario $s$ depends on the value $s(p)$ only for those variables $p \in \mathcal{P}$ that are observed strictly before the time point $[\sigma(s)]_N$. This condition is sufficient to guarantee that $\sigma$ is dynamic. In particular, observe that the condition $i \leq b(s)$ depends only on the variables $x_1, y_1, \ldots, x_{i-1}, y_{i-1}$, and that $[\sigma(s)]_{D_i}$ for $i \leq b'(s)$ depends on either $c_i^0$ or $c_i^1$, whichever has been actually observed at time $(n+4)i$ in the scenario $s$.

One can easily check that $\sigma$ is viable, by checking that all the constraints in $\mathcal{C}$ are satisfied. In particular, concerning the constraints $(B_{n+1} \geq A_{n+1} + (n+1), \neg l_{j,1} \wedge \neg l_{j,2} \wedge \neg l_{j,3})$ for $j \in \{1, \ldots, m\}$, there are two possibilities. If $b(s) =$

$$
\begin{array}{lll}
[\sigma(s)]_{A_i} = (n+4)\,i & \text{for } i = 1,\dots,n+1 \\
[\sigma(s)]_{B_i} = (n+4)\,i & \text{for } i = 1,\dots,b(s) \\
[\sigma(s)]_{B_i} = \infty & \text{for } i = b(s)+1,\dots,n+1 \\
[\sigma(s)]_{C_i^h} = (n+4)\,i & \text{for } i = 1,\dots,b'(s) \text{ and } h = h_i(s) \\
[\sigma(s)]_{C_i^h} = \infty & \text{for } i = 1,\dots,b'(s) \text{ and } h \neq h_i(s) \\
[\sigma(s)]_{C_i^h} = \infty & \text{for } i = b'(s)+1,\dots,n \text{ and } h = 0,1 \\
[\sigma(s)]_{D_i} = \begin{cases} (n+4)\,i+1 & \text{if } s(c_i^{h_i(s)}) = 1 \\ \infty & \text{otherwise} \end{cases} & \\
\qquad\qquad \text{for } i = 1,\dots,b'(s) & \\
[\sigma(s)]_{D_i} = \infty & \text{for } i = b'(s)+1,\dots,n \\
[\sigma(s)]_{X_i} = (n+4)\,i+n+2 & \text{for } i = 1,\dots,n+1 \\
[\sigma(s)]_{Y_i} = (n+4)\,i+n+3 & \text{for } i = 1,\dots,n+1.
\end{array}
$$

Figure 3. Dynamic and viable execution strategy $\sigma$ for $\Gamma_\Phi$, when $\Phi \equiv$ true. Fixed a winning strategy $f$ for $\Phi$, the execution strategy $\sigma$ is defined above, where $h_i(s) \coloneqq f(s(y_1),\dots,s(y_{i-1}))$, $b(s) \coloneqq \min\{i \mid s(x_i) \neq h_i(s)\} \cup \{n+1\}$ and $b'(s) \coloneqq \min\{b(s), n\}$.

$n+1$, then all these constraint are void since each clause $(l_{j,1} \vee l_{j,2} \vee l_{j,3})$ of $\varphi$ is satisfied by the interpretation given by $s$. Otherwise, if $b(s) \in \{1,\dots,n\}$, then they are all satisfied since $[\sigma(s)]_{B_{n+1}} = \infty \geq [\sigma(s)]_{A_{n+1}} + (n+1) = (n+4)(n+1) + n + 1$. ∎

To prove the converse of Lemma 2, we first spot out three facts detailing out how the gadgets work as intended. First, if the $i$-the gadget is activated, then the controller is forced to execute either $C_i^0$ or $C_i^1$ early (Lemma 3). Second, the activation of the $i$-th gadget and the choice of either $C_i^0$ or $C_i^1$ cannot depend on the variables $x_i, y_i, \dots, x_n, y_n$, so the nature can choose their values "later on" (Lemma 4). Third, if the nature copies the value selected by the controller, the activation constraint is propagated to the next gadget (Lemma 5).

**Lemma 3** (The controller has to schedule $C_i^0$ or $C_i^1$ early)**.** *Let $\sigma$ be a viable and dynamic execution for $\Gamma_\Phi$. Let $s$ be any scenario and $i \in \{1,\dots,n\}$, and suppose that $[\sigma(s)]_{B_i} - [\sigma(s)]_{A_i} \leq n$. Then, for some $h \in \{0,1\}$, we have $[\sigma(s)]_{C_i^h} \leq [\sigma(s)]_{B_i} + 1$.*

*Proof:* Fix a scenario $s$, let $t \coloneqq [\sigma(s)]_{B_i} + 1$, and suppose by contradiction $[\sigma(s)]_{C_i^h} \geq t$ for $h = 0, 1$. Let $d = 1$ if $[\sigma(s)]_{D_i} \geq t$ and $d = 0$ otherwise, and take $s' = s[d/c_i^0][d/c_i^1]$. Since $t \leq [\sigma(s)]_{C_i^h}$ for $h = 0, 1$, we can apply Lemma 1 obtaining that $[\sigma(s')]_{B_i} + 1 = t$, $[\sigma(s')]_{C_i^h} \geq t$, $[\sigma(s')]_{D_i} \geq t \iff [\sigma(s)]_{D_i} \geq t \iff d = 0$, and either $[\sigma(s)]_{A_i} = [\sigma(s')]_{A_i} < t$ or both $[\sigma(s)]_{A_i}, [\sigma(s')]_{A_i} \geq t$. Now, if $d = 1$, then the constraint $D_i \leq B_i + 1$ applies in scenario $s'$ and is violated by $\sigma(s')$. Otherwise, if $d = 0$, the constraint $D_i \geq A_i + n + 1$ applies in scenario $s'$ and is violated by $\sigma(s')$. In either case, this contradicts the fact that $\sigma$ is viable. ∎

**Lemma 4** (The nature can choose future variables)**.** *Let $\sigma$ be a viable and dynamic execution for $\Gamma_\Phi$. Let $s$ be any scenario and $i \in \{1,\dots,n\}$. Let $s' = s[v/p]$ be a scenario obtained by changing the value of any variable $p \in \{x_i, y_i, \dots, x_n, y_n\}$ to any value $v \in \{0,1\}$. Then, we have $[\sigma(s')]_{B_i} - [\sigma(s')]_{A_i} \leq i-1 \iff [\sigma(s)]_{B_i} - [\sigma(s)]_{A_i} \leq i-1$. Moreover, if $[\sigma(s)]_{B_i} - [\sigma(s)]_{A_i} \leq i-1$ holds, then $[\sigma(s')]_{C_i^h} < [\sigma(s')]_{B_i} + 1 \iff [\sigma(s)]_{C_i^h} < [\sigma(s)]_{B_i} + 1$ for both $h = 0$ and $h = 1$.*

*Proof:* Let $t = [\sigma(s)]_{A_i} + n + 1$. Since $\sigma$ is viable, the unlabeled constraints $X_i \geq A_i + (n+2)$, $Y_i \geq X_i + 1$ and $A_{i+1} \geq Y_i + 1$ imply that $t = [\sigma(s)]_{A_i} + n + 1 \leq [\sigma(s)]_{X_i} \leq [\sigma(s)]_{Y_i} \leq \cdots \leq [\sigma(s)]_{X_n} \leq [\sigma(s)]_{Y_n}$. Since $[\sigma(s)]_{A_i} \leq [\sigma(s)]_{A_i} + i - 1 \leq t$, Lemma 1 can be applied to obtain that $[\sigma(s')]_{A_i} = [\sigma(s)]_{A_i}$ and $[\sigma(s)]_{B_i} \leq [\sigma(s)]_{A_i} + i - 1 \iff [\sigma(s')]_{B_i} \leq [\sigma(s')]_{A_i} + i - 1$ which proves the first part of the lemma. Now assume $[\sigma(s')]_{B_i} \leq [\sigma(s)]_{A_i} + i - 1$: we have

$$
\begin{array}{rcl}
[\sigma(s')]_{B_i} \leq [\sigma(s')]_{B_i} + 1 & \leq & [\sigma(s')]_{A_i} + (i-1) + 1 \\
& \leq & [\sigma(s')]_{A_i} + n + 1 = t
\end{array}
$$

so by Lemma 1 we obtain that $[\sigma(s')]_{B_i} = [\sigma(s)]_{B_i}$ and $[\sigma(s')]_{C_i^h} < [\sigma(s')]_{B_i} + 1 \iff [\sigma(s)]_{C_i^h} < [\sigma(s)]_{B_i} + 1$. ∎

**Lemma 5** (Propagation of the activation constraint)**.** *Let $\sigma$ be a viable and dynamic execution for $\Gamma_\Phi$. Let $s$ be any scenario and $i \in \{1,\dots,n\}$, and suppose that $[\sigma(s)]_{B_i} - [\sigma(s)]_{A_i} \leq i-1$. Let $h = 0$ if $[\sigma(s)]_{C_i^0} < [\sigma(s)]_{B_i} + 1$ and $h = 1$ otherwise. For $s' = s[h/x_i]$ we have $[\sigma(s')]_{B_{i+1}} - [\sigma(s')]_{A_{i+1}} \leq i$.*

*Proof:* If $h = 1$ then $[\sigma(s)]_{C_i^1} < [\sigma(s)]_{B_i} + 1$ by Lemma 3. From Lemma 4 we obtain that $[\sigma(s')]_{B_i} - [\sigma(s')]_{A_i} \leq i-1$ and $[\sigma(s')]_{C_i^h} < [\sigma(s')]_{B_i} + 1$. Moreover, since $\sigma$ is viable, thanks to the constraint $B_{i+1} \leq C_i^h + n + 4$, labeled $x_i$ if $h = 1$ and $\neg x_i$ otherwise, we obtain that $[\sigma(s)]_{B_{i+1}} \leq [\sigma(s)]_{C_i^h} + n + 4 < [\sigma(s)]_{B_i} + 1 + n + 4 \leq [\sigma(s)]_{A_i} + i - 1 + 1 + n + 4 \leq [\sigma(s)]_{A_{i+1}} + i$ where the last inequality follows from the unlabeled constraints $A_{i+1} \geq Y_i + 1$, $Y_i \geq X_i + 1$ and $X_i \geq A_i + (n+2)$. ∎

**Lemma 6.** *If $\Phi \equiv$ false then $\Gamma_\Phi$ is not dynamically controllable.*

*Proof:* Let $f\colon \{0,1\}^* \to \{0,1\}$ be the winning strategy of the universal player for $\Phi$. Suppose by contradiction that $\sigma$ is a viable and dynamic execution strategy for $\Gamma_\Phi$.

We first construct, for $I = 0,\dots,n$, step by step, a scenario $s_I$ such that

(a) $[\sigma(s_I)]_{B_{I+1}} - [\sigma(s_I)]_{A_{I+1}} \leq I$ (activation constraint), *and*

(b) $s_I(y_i) = f(s(x_1),\dots,s(x_i))$ for $i = 1,\dots,I$.

Start with any scenario $s_0$. We have (a) $[\sigma(s_0)]_{B_1} - [\sigma(s_I)]_{A_1} \leq 0$ thanks to the constraint $B_1 - A_1 \leq 0$, and there is nothing to prove for (b). For $I = 0,\dots,n-1$,

define $s_{I+1}$ as follows. Let $h_{I+1} = 0$ if $[\sigma(s_I)]_{C^0_{I+1}} < [\sigma(s_I)]_{B_{I+1}} + 1$ and $h_{I+1} = 1$ otherwise, and define $s_{I+1} = s_I[h_{I+1}/x_{I+1}][f(s_I(x_1), \ldots, s_I(x_I), h_{I+1})/y_{I+1}]$. By construction (b) is satisfied. We obtain (a) by applying Lemma 4 and Lemma 5.

Consider the scenario $s_n$. We have (a) $[\sigma(s_n)]_{B_{n+1}} - [\sigma(s_n)]_{A_{n+1}} \leq n$. Moreover, by (b) and the fact that $f$ is a winning strategy for the universal player, the formula $\varphi$ is false in the interpretation given by the scenario $s_n$. In particular, some clause is not satisfied, say, the $j$-th clause for some $j \in \{1, \ldots m\}$. So, the constraint $B_{n+1} - A_{n+1} \geq n+1$ labeled with $\ell_j$ applies in scenario $s_n$, but it is violated since we proved $[\sigma(s_n)]_{B_{n+1}} - [\sigma(s_n)]_{A_{n+1}} \leq n$. ■

**Theorem 1.** *CSTN-DC is PSPACE-hard.*

*Proof:* Given a Q3SAT formula $\Phi$, the CSTN $\Gamma_\Phi$ can be easily constructed within logarithmic internal memory. By Lemmas 2 and 6, it is dynamically controllable iff $\Phi \equiv$ true. ■

## IV. Polynomial-space algorithm

### A. Relative execution strategies

First, we extend some of the notions for CSTNs to the case when some of the tasks have already been performed. This will be crucial to describe our inductive polynomial-space algorithm.

**Definition 15** (Partial schedule, next action, completion)**.** A *partial schedule* over $\mathcal{T}$ up to time $t \in \mathbb{R}$ is a schedule $\psi$ over a subset $\mathrm{Dom}(\psi) \subseteq \mathcal{T}$, such that $[\psi]_X \leq t$ for every $X \in \mathrm{Dom}(\psi)$. Given a partial schedule $\psi$ up to time $t$, a *next action* for $\psi$ is a pair $(t_{next}, \mathcal{T}_{next})$ where $t_{next} > t$ is a time point and $\mathcal{T}_{next} \subseteq \mathcal{T} \setminus \mathrm{Dom}(\psi)$ is a non-empty set of temporal variables not assigned by $\psi$. Let $\psi[t_{next}/\mathcal{T}_{next}] = \psi \cup \{(X, t_{next}) \mid X \in \mathcal{T}_{next}\}$ be the partial schedule, up to time $t_{next}$, obtained from $\psi$ by further executing all the actions in $\mathcal{T}_{next}$ at time $t_{next}$. Given a partial schedule $\psi$ up to time $t$, a *completion* of $\psi$ is a schedule $\psi' \in \Psi_\mathcal{T}$ such that $[\psi']_X = [\psi]_X$ for every $X \in \mathrm{Dom}(\psi)$ and $[\psi']_X > t$ for every $X \in \mathrm{Dom}(\psi') \setminus \mathrm{Dom}(\psi)$. Let $\Psi_\mathcal{T}[\psi]$ be the set of completions of $\psi$.

**Definition 16** (Observation, completion of a partial scenario)**.** Given a partial scenario $h$ and a set of propositional variables $\mathcal{P}' \subseteq \mathcal{P} \setminus \mathrm{Dom}(h)$ not assigned by $h$, an *observation* of $\mathcal{P}'$ is a function $o\colon \mathcal{P}' \to \{0,1\}$, and $h \cup o$ is the partial scenario obtained from $h$ by adding all the assignments given by $o$. Given a partial scenario $h$, a *completion* of $h$ is a total scenario $s \in \Sigma_\mathcal{P}$ such that $s(p) = h(p)$ for every $p \in \mathrm{Dom}(h)$. Let $\Sigma_\mathcal{P}[h]$ denote the set of completions of $h$.

**Definition 17** (Configuration, initial, terminal)**.** A *configuration* is a tuple $c = (t, \psi, h)$ consisting of a time point $t \in \mathbb{R} \cup \{-\infty\}$, a partial schedule $\psi$ up to time $t$, and a partial scenario $h\colon \mathcal{P}_c \to \{0,1\}$ where $\mathcal{P}_c = \{p \in \mathcal{P} \mid \mathcal{O}(p) \in \mathrm{Dom}(\psi)\}$ is the set of propositional variables observed before or at time $t$. Let $c_0 = (-\infty, \psi_0, h_0)$ be the *initial configuration*, where $\mathrm{Dom}(\psi_0) = \emptyset$ and $\mathrm{Dom}(h_0) = \emptyset$. A configuration $c = (t, \psi, h)$ is *terminal* if, for every scenario $s \in \Sigma_\mathcal{P}[h]$, we have $\mathcal{T}_s = \mathrm{Dom}(\psi)$.

**Definition 18** (Next configuration)**.** Given a configuration $c = (t, \psi, h)$, a next action $(t_{next}, \mathcal{T}_{next})$ for $\psi$, and an observation $o\colon \mathcal{P}_{next} \to \{0,1\}$ of $\mathcal{P}_{next} \coloneqq \{p \in \mathcal{P} \mid \mathcal{O}(p) \in \mathcal{T}_{next}\}$, define the *next configuration* $c[t_{next}/\mathcal{T}_{next}, o] = (t_{next}, \psi[t_{next}/\mathcal{T}_{next}], h \cup o)$.

**Definition 19** (Relative execution strategies)**.** A *relative execution strategy* from a configuration $c = (t, \psi, h)$ is a function $\sigma\colon \Sigma_\mathcal{P}[h] \to \Psi_\mathcal{T}[\psi]$ that maps each scenario $s$ which is a completion of $h$ to a total schedule $\sigma(s)$, over $\mathcal{T}_s$, which is a completion of $\psi$. A relative execution strategy $\sigma$ is *viable* if $\sigma(s)$ is feasible for $\Gamma_s$ for every scenario $s \in \Sigma_\mathcal{P}[h]$. It is *dynamic* if, for any scenarios $s, s' \in \Sigma_\mathcal{P}[h]$ and time variable $X \in \mathcal{T}_s$, letting $t = [\sigma(s)]_X$, if $\mathrm{Hist}(t, s, \sigma) = \mathrm{Hist}(t, s', \sigma)$ then $X \in \mathcal{T}_{s'}$ and $[\sigma(s')]_X = t$. Observe that a (dynamic, viable) relative execution strategy from the initial configuration $c_0$ is a (dynamic, viable) execution strategy and vice-versa.

**Definition 20** (Dynamic controllability from a configuration)**.** A CSTN is *dynamically controllable* from a configuration $c$ if it admits a dynamic and viable relative execution strategy from $c$.

The following definition and lemma serve to ensure that, in a dynamic relative execution strategy from a non-terminal configuration, there is always a set of actions that is executed next, all at the same time across all the scenarios.

**Definition 21** (Well-defined next action)**.** A relative execution strategy $\sigma$ from a configuration $c = (t, \psi, h)$ has a *well-defined next action* if there exists a next action $(t_{next}(\sigma), \mathcal{T}_{next}(\sigma))$ for $\psi$ such that, for every scenario $s \in \Sigma_\mathcal{P}[h]$, $\min_{X \in \mathcal{T}_s \setminus \mathrm{Dom}(\psi)} [\sigma(s)]_X = t_{next}(\sigma)$ and $\{X \in \mathcal{T}_s \mid [\sigma(s)]_X = t_{next}(\sigma)\} = \mathcal{T}_{next}(\sigma)$. Equivalently, $\sigma$ has a well-defined next action if, for every scenario $s \in \Sigma_\mathcal{P}[h]$, $\sigma(s)$ is a completion of $\psi[t_{next}(\sigma)/\mathcal{T}_{next}(\sigma)]$.

**Lemma 7.** *If $\sigma$ is a dynamic execution strategy from a non-terminal configuration $c = (t, \psi, h)$, then $\sigma$ has a well-defined next action.*

*Proof:* Since $\sigma$ is non-terminal there exist some $s \in \Sigma_\mathcal{P}[h]$ and $X \in \mathcal{T}_s \setminus \mathrm{Dom}(\psi)$. Therefore, we can define $t_{next}(\sigma) = \min_{s \in \Sigma_\mathcal{P}[h], X \in \mathcal{T}_s \setminus \mathrm{Dom}(\psi)} [\sigma(s)]_X$. Then, we choose any $s_0 \in \Sigma_\mathcal{P}[h]$ and define $\mathcal{T}_{next}(\sigma) = \{X \in \mathcal{T}_{s_0} \mid [\sigma(s_0)]_X = t_{next}(\sigma)\}$. We prove that $\mathcal{T}_{next}(\sigma)$ does not depend on the choice of $s_0$, using the fact that $\sigma$ is dynamic, and this concludes the proof.

Take any $s, s' \in \Sigma_\mathcal{P}[h]$, and suppose $[\sigma(s)]_X = t_{next}(\sigma)$.

We want to prove that also $[\sigma(s')]_X = t_{next}(\sigma)$. By definition of $t_{next}(\sigma)$, there is no $X' \in \mathcal{T} \setminus \mathrm{Dom}(\psi)$ with either $[\sigma(s)]_X < t_{next}(\sigma)$ or $[\sigma(s')]_X < t_{next}(\sigma)$. On the other hand, for every $X' \in \mathrm{Dom}(\psi)$, we have both $[\sigma(s)]_X < t_{next}(\sigma)$ and $[\sigma(s')]_X < t_{next}(\sigma)$. Hence, $\mathrm{Hist}(t_{next}(\sigma), s, \sigma) = \mathrm{Hist}(t_{next}(\sigma), s', \sigma) = h$. Since $\sigma$ is dynamic, by applying the definition we obtain $[\sigma(s')]_X = t_{next}(\sigma)$ as desired. ■

**Definition 22** (Child configurations and strategies)**.** Suppose $\sigma$ has a well-defined next action. Let $\mathcal{P}_{next}(\sigma) = \{p \in \mathcal{P} \mid \mathcal{O}(p) \in \mathcal{T}_{next}(\sigma)\}$ be the set of propositional variables observed at time $t_{next}(\sigma)$ and $o\colon \mathcal{P}_{next}(\sigma) \to \{0,1\}$ be any outcome for the observations. Define the *child configuration* $next(\sigma, o) = c[t_{next}(\sigma)/\mathcal{T}_{next}(\sigma), o] = (t_{next}(\sigma), \psi[t_{next}(\sigma)/\mathcal{T}_{next}(\sigma)], h \cup o)$. Since, for every $s \in \Sigma_{\mathcal{P}}[h]$, the schedule $\sigma(s)$ is a completion of $\psi[t_{next}(\sigma)/\mathcal{T}_{next}(\sigma)]$, a relative execution strategy from $next(\sigma, o)$ is obtained simply restricting $\sigma$ to the scenarios that are completions of $h \cup o$. Denote this strategy with $child(\sigma, o) := \sigma|_{\Sigma_{\mathcal{P}}[h \cup o]}$.

**Lemma 8.** *If $\sigma$ is dynamic then also $child(\sigma, o)$ is dynamic.*

*Proof:* Since $child(\sigma, o)$ is a restriction of $\sigma$, there are less pairs $s, s' \in \Sigma_{\mathcal{P}}$ that need to be checked in order for $child(\sigma, o)$ to be dynamic. ■

**Lemma 9.** *Let $\sigma$ be a relative execution strategy from a non-terminal configuration $c = (t, \psi, h)$. If $\sigma$ has a well-defined next action and $child(\sigma, o)$ is dynamic for every $o\colon \mathcal{P}_{next}(\sigma) \to \{0,1\}$, then $\sigma$ is dynamic.*

*Proof:* Let $s, s' \in \Sigma_{\mathcal{P}}[h]$, $X \in \mathcal{T}_s$, with $t = [\sigma(s)]_X$ and $\mathrm{Hist}(t, s, \sigma) = \mathrm{Hist}(t, s', \sigma)$. We need to prove that $[\sigma(s')]_X = t$. If $t < t_{next}(\sigma)$, then $X \in \mathrm{Dom}(\psi)$ so $[\sigma(s)]_X = [\psi]_X = [\sigma(s')]_X = t$. If $t = t_{next}(\sigma)$, then $X \in \mathcal{T}_{next}(\sigma)$ so $[\sigma(s')]_X = t_{next}(\sigma) = [\sigma(s)]_X = t$. If $t > t_{next}(\sigma)$, then for any $p \in \mathcal{P}_{next}(\sigma)$, we have $[\sigma(s)]_{\mathcal{O}(p)} = t_{next}(\sigma) < t$. So, $p \in \mathrm{Dom}(\mathrm{Hist}(t, s, \sigma))$ and, since $\mathrm{Hist}(t, s, \sigma) = \mathrm{Hist}(t, s', \sigma)$ by hypothesis, we have $s(p) = s'(p)$. Take $o\colon \mathcal{P}_{next}(\sigma) \to \{0,1\}$ so that $o(p) = s(p) = s'(p)$ for every $p \in \mathcal{P}_{next}(\sigma)$. Since $child(\sigma, o)$ is dynamic by hypothesis, and both $s, s' \in \Sigma_{\mathcal{P}}[h \cup o]$, by the definition of dynamic strategy we get $[\sigma(s')]_X = t$. ■

**Lemma 10.** *Let $\sigma$ be a dynamic relative execution strategy from a non-terminal configuration $c$. Then, $\sigma$ is viable iff $child(\sigma, o)$ is viable for every $o\colon \mathcal{P}_{next}(\sigma) \to \{0,1\}$.*

*Proof:* $\sigma = \bigcup_{o\colon \mathcal{P}_{next}(\sigma) \to \{0,1\}} child(\sigma, o)$. ■

**Lemma 11.** *A CSTN $\Gamma$ is dynamically controllable from a* terminal *configuration $c = (t, \psi, h)$ iff, for every scenario $s \in \Sigma_{\mathcal{P}}[h]$, the schedule $\psi$ is feasible for $\Gamma_s$.*

*Proof:* There exists only one execution strategy $\sigma$ from $c$, defined by $\sigma(s) = \psi$ for every $s \in \Sigma_{\mathcal{P}}[h]$. It is clearly dynamic, and, by definition, it is viable iff, for every scenario $s \in \Sigma_{\mathcal{P}}[h]$, the schedule $\psi$ is feasible for $\Gamma_s$ ■

**Lemma 12.** *A CSTN $\Gamma$ is dynamically controllable from a* non-terminal *configuration $c = (t, \psi, h)$ iff there exist a next action $(\mathcal{T}_{next}, t_{next})$ from $\psi$ such that, for every observation $o\colon \mathcal{P}_{next} \to \{0,1\}$ (where $\mathcal{P}_{next} = \{p \in \mathcal{P} \mid \mathcal{O}(p) \in \mathcal{T}_{next}\}$), $\Gamma$ is dynamically controllable from $c[t_{next}/\mathcal{T}_{next}, o]$.*

*Proof:* ( $\Longrightarrow$ ) Let $\sigma$ be a dynamic and viable execution strategy from $c$. It is sufficient to apply Lemma 7 and take $t_{next} = t_{next}(\sigma)$ and $\mathcal{T}_{next} = \mathcal{T}_{next}(\sigma)$. Then, for every $o\colon \mathcal{P}_{next} \to \{0,1\}$, the strategy $child(\sigma, o)$ is dynamic and viable from $c[t_{next}/\mathcal{T}_{next}, o]$, thanks to Lemma 8 and Lemma 10.

( $\Longleftarrow$ ) For every $o\colon \mathcal{P}_{next} \to \{0,1\}$, let $\sigma_o$ be a dynamic and viable execution strategy from $c[t_{next}/\mathcal{T}_{next}, o]$. Then, define the strategy $\sigma = \bigcup_{o\colon \mathcal{P}_{next} \to \{0,1\}} \sigma_o$ from the configuration $c$. Observe that $\sigma$ has a well-defined next action, and in particular $t_{next}(\sigma) = t_{next}$ and $\mathcal{T}_{next}(\sigma) = \mathcal{T}_{next}$. Moreover, for every $o\colon \mathcal{P}_{next} \to \{0,1\}$, we have $child(\sigma, o) = \sigma_o$ which is dynamic and viable by assumption. Thanks to Lemma 9 and Lemma 10, $\sigma$ is dynamic and viable. ■

Lemma 11 (base case) and Lemma 12 (inductive case) suggest a recursive approach to solve CSTN-DC. In the inductive case, we need to consider all the possible choices of $t_{next}$ and $\mathcal{T}_{next}$. However, this is still not possible since $t_{next}$ is, a priori, an unbounded real number. In the following we show that, under suitable assumptions, we can choose $t_{next}$ among a finite set of possibilities.

### *B. Discrete strategies for CSTNs*

We assume to work on CSTNs whose constraint bounds are discrete, and can be expressed with a finite number of bits in fixed-point precision. This is stated in the following definition.

**Definition 23** (Discrete CSTN)**.** Let $w \in \mathbb{R}$ and $W \in \mathbb{N}$. A CSTN is $(w, W)$-discrete if, for every labeled constraint $(Y \leq X + \delta, \ell) \in \mathcal{C}$, we have $\delta = kw$ for $k \in \{-W, \ldots, +W\}$. We call *Discrete CSTN-DC* the variant of CSTN-DC where the input CSTN is $(w, W)$-discrete for some $w \in \mathbb{R}$ and $W \in \mathbb{N}$.

We prove that, for discrete CSTNs, one can always restrict her attention to discrete execution strategies, whose execution times are expressible with a number of bits at most polynomial in the size of the input. Our proof is a generalization of an argument given in [5].

**Definition 24** (Discrete execution strategy)**.** Let $\mu \in \mathbb{R}$ and $M \in \mathbb{N}$. A (relative) execution strategy $\sigma$ is $(\mu, M)$*-discrete* if $[\sigma(s)] = k\mu$, with $k \in \{1, \ldots, M\}$, for every scenario $s \in \Sigma_{\mathcal{P}}$ and time variable $X \in \mathcal{T}_s$.

**Lemma 13** (Discrete CSTNs admit discrete strategies)**.** *Consider a $(w, W)$-discrete CSTN $\Gamma$. If $\Gamma$ is dynamically*

*controllable, then* $\Gamma$ *admits a* $(\mu, M)$*-discrete viable dynamic execution strategy, for* $\mu \coloneqq w/K$*,* $M \coloneqq 2 \cdot K^2 \cdot W$ *and* $K \coloneqq 2^{|\mathcal{P}|} \cdot |\mathcal{T}|$*.*

*Proof:* Let $\sigma$ be a viable dynamic strategy for $\Gamma$. For each $s \in \Sigma_{\mathcal{P}}$ and $X \in \mathcal{T}_s$, write $[\sigma(s)]_X$ as

$$[\sigma(s)]_X = a_{s,X} \cdot W \cdot w + b_{s,X} \cdot w + c_{s,X}$$

for $a_{s,X} \in \mathbb{Z}$, $b_{s,X} \in \{0, \dots, W-1\}$ and $c_{s,X} \in [0, w)$. Let $A = \{a_{s,X},\, a_{s,X}+1 \mid s \in \Sigma_{\mathcal{P}}, X \in \mathcal{T}_s\}$ and $C = \{c_{s,X} \mid s \in \Sigma_{\mathcal{P}}, X \in \mathcal{T}_s\}$. Then, let $\alpha_{s,X} \in \{0, \dots, 2K-1\}$ be the 0-based rank respectively of $a_{s,X}$ in $A$ and $\gamma_{s,X} \in \{0, \dots, K-1\}$ the 1-based rank of $c_{s,X}$ in $C$. Define the strategy $\sigma'$ as follows

$$[\sigma'(s)]_X = \alpha_{s,X} \cdot W \cdot w + b_{s,X} \cdot w + \gamma_{s,X} \cdot w/K.$$

By construction, $\sigma'$ is $(\mu, M)$-discrete for $\mu \coloneqq w/K$ and $M \coloneqq 2 \cdot K^2 \cdot W$. We show that $\sigma'$ is viable and dynamic, thus proving the statement. Observe that $[\sigma'(s)]_X < [\sigma'(s')]_Y \iff [\sigma(s)]_X < [\sigma(s')]_Y$, for every $s \in \Sigma_{\mathcal{P}}$ and $X \in \mathcal{T}_s$.

(Viable.) Let $(Y \le X + kw, \ell)$ be a constraint of $\Gamma$ and $s \in \Sigma_{\mathcal{P}}$ a scenario such that $s \vDash \ell$. We have $[\sigma(s)]_Y - [\sigma(s)]_X \le kw$ by the assumption that $\sigma$ is viable. We prove that $[\sigma'(s)]_Y - [\sigma'(s)]_X \le kw$, distinguishing among the following cases.

1) Case $|a_{s,Y} - a_{s,X}| \ge 2$.
   Then also $|\alpha_{s,Y} - \alpha_{s,X}| \ge 2$, since we added both $\alpha_{s,Z}$ and $\alpha_{s,Z}+1$ to $A$, for $Z \in \{X, Y\}$. Hence, $|[\sigma'(s)]_Y - [\sigma'(s)]_X| > W$ and, since $|k| \le W$, we have $[\sigma'(s)]_Y - [\sigma'(s)]_X \le kw$.
2) Case $a_{s,Y} - a_{s,X} \in \{-1, 0, +1\}$.
   We have $a_{s,Y} - a_{s,X} = \alpha_{s,Y} - \alpha_{s,X}$.
   a) Case $(a_{s,Y} \cdot W + b_{s,Y}) - (a_{s,X} \cdot W + b_{s,X}) \le k-1$.
      Then, $[\sigma'(s)]_Y - [\sigma'(s)]_X < (k-1) \cdot w + w \le kw$.
   b) Case $(a_{s,Y} \cdot W + b_{s,Y}) - (a_{s,X} \cdot W + b_{s,X}) = k$.
      Then $c_{s,Y} \le c_{s,X}$, so also $\gamma_{s,Y} \le \gamma_{s,X}$ and $[\sigma'(s)]_Y - [\sigma'(s)]_X \le kw + (\gamma_{s,Y} - \gamma_{s,X}) \le kw$.

(Dynamic.) The fact that an execution strategy is dynamic depends only on the relative order (in $\mathbb{R}$) of the values of $[\sigma(s)]_X$, which is preserved by our transformation to $\sigma'$. ■

### C. The algorithm

We first adapt Lemma 12 to relative execution strategies.

**Definition 25** (Discrete configuration)**.** A configuration $c = (t, \psi, h)$ is $(\mu, M)$-discrete if $[\psi]_X = k\mu$, with $k \in \{1, \dots, M\}$, for every time variable $X \in \mathrm{Dom}(\psi)$.

**Lemma 14.** *A CSTN* $\Gamma$ *admits a* $(\mu, M)$*-discrete dynamic and viable execution strategy from a* non-terminal *$(\mu, M)$-discrete configuration* $c = (t, \psi, h)$*, iff there exist a next action* $(\mathcal{T}_{next}, t_{next})$ *from* $\psi$*, with* $t_{next} = k_{next}\,\mu > t$ *and* $k_{next} \in \{1, \dots, M\}$*, such that, for every observation* $o\colon \mathcal{P}_{next} \to \{0, 1\}$ *(where* $\mathcal{P}_{next} = \{p \in \mathcal{P} \mid \mathcal{O}(p) \in \mathcal{T}_{next}\}$*),* $\Gamma$ *admits a* $(\mu, M)$*-discrete dynamic and viable execution strategy from the* $(\mu, M)$*-discrete configuration* $c[t_{next}/\mathcal{T}_{next}, o]$*.*

*Proof:* Trivial adaptation of the proof of Lemma 12. ■

The polynomial-memory algorithm is a mere application of Lemma 13, Lemma 11 and Lemma 14. The pseudo-code is shown in Algorithm 1.

**Lemma 15.** *Algorithm 1 can be implemented using at most* $O(|\mathcal{T}| \cdot (\log |W| + \log |\mathcal{T}|))$ *space.*

*Proof:* To implement the recursive procedure, it is sufficient to maintain a stack of triples $(\mathcal{T}_1, k_1, o_1) \cdots (\mathcal{T}_l, k_l, o_l)$ containing the choices of $\mathcal{T}_{next}$, $k_{next}$ and $o$ at each level of the recursion. Indeed, the parameters $\psi$ and $h$ can be reconstructed from this stack. To represent the sequence $k_1, \dots, k_l$, we need $O(|\mathcal{T}| \log |M|) = O(|\mathcal{T}|(\log |W| + \log |\mathcal{T}|))$ bits, while $\mathcal{T}_1, \dots, \mathcal{T}_l$ and $o_1, \dots, o_l$ require only $O(|\mathcal{T}| \log |\mathcal{T}|)$ bits. ■

As a consequence we get the following.

**Theorem 2.** *Discrete CSTN-DC is in PSPACE.*

### D. Extending to real-valued CSTNs

It is possible to extend our polynomial algorithm so that it works without making any assumption on how the input numbers are encoded. Since Lemma 13 does not apply, we need a different way to limit the choice of $t_{next}$ to a finite set. We verified that it is sufficient to take $t_{next}$ among those linear combinations of the input numbers, having integer coefficients with a number of bits polynomial in the size of the network. A formal proof of this statement, as well as the extension of this positive result to CSTNUs, is subject of future work.

## V. Conclusion

Our first result is a reduction from Q3SAT to CSTN-DC, which shows that checking the dynamic controllability of CSTNs is PSPACE-hard. Our reduction relies on the close interplay between labeled constrains and observation tasks, which allows the planner to effectively impose her choice on some of the propositional variables. This interplay seems to be the reason why CSTN-DC is difficult. On the other hand, the topology of the network plays very little role. Indeed, in our construction, the topology of the network is planar and extremely simple: there is only one directed cycle (observe that, if there was no directed cycle, then it would be trivially and strongly controllable), and removing a *single* edge, from $A_1$ to $B_1$, we get an acyclic graph, and actually an inward arborescence (a directed rooted tree, where all edges point towards the root) when disregarding the parallel edges between $A_n$ and $B_n$.

Our second result is an algorithm for CSTN-DC that uses only polynomial space, proving that CSTN-DC $\in$ PSPACE.

**Algorithm 1:** *Discrete CSTN-DC* in polynomial space.

**Function** *DC*$(\Gamma)$
  **Input** : $\Gamma$ is a $(w, W)$-discrete CSTN
  **Returns**: *true* if $\Gamma$ is dynamic controllable, *false* otherwise
  $c_0 \coloneqq (0, \emptyset, \emptyset)$ ▷ Initial configuration
  **return** *DC-From*$(\Gamma, c_0)$

**Recursive Function** *DC-From*$(\Gamma, c)$
  **Input** : $\Gamma$ is a $(w, W)$-discrete CSTN
    $c = (k\mu, \psi, h)$ is a $(\mu, M)$-discrete configuration, for $\mu \coloneqq w/K$, $M \coloneqq 2\,K^2 \cdot W$ and $K \coloneqq 2^{|\mathcal{P}|} \cdot |\mathcal{T}|$.
  **Returns**: *true* if $\Gamma$ is dynamic controllable from $c$, *false* otherwise
  **if** *Terminal-And-DC*$(\Gamma, c)$ **then**
    **return** *true*
  ▷ Enumerate all the possible next actions ($\exists$-step)
  **foreach** $\mathcal{T}_{next} \subseteq \mathcal{T} \setminus \mathrm{Dom}(\psi)$ *not empty* **and**
  **foreach** $k_{next} \in \{k+1, \ldots, M\}$ **do**
    $t_{next} \coloneqq k_{next}\,\mu$
    $\psi' \coloneqq \psi[t_{next}/\mathcal{T}_{next}]$
    $\mathcal{P}_{next} \coloneqq \{p \in \mathcal{P} \mid \mathcal{O}(p) \in \mathcal{T}_{next}\}$
    ▷ Enumerate all the possible observations ($\forall$-step)
    $AllChildrenAreDC \leftarrow$ true
    **foreach** $o \colon \mathcal{P}_{next} \to \{0, 1\}$ **do**
      $h' \coloneqq h \cup o$
      $c' \coloneqq (t_{next}, \psi', h')$
      **if** *not DC-From*$(\Gamma, c')$ **then** ▷ Recursion
        $AllChildrenAreDC \leftarrow$ false
    **if** $AllChildrenAreDC$ **then**
      **return** *true*
  **return** *false*

**Function** *Terminal-And-DC*$(\Gamma, c)$
  **Input** : CSTN $\Gamma$, and configuration $c = (t, \psi, h)$
  **Returns**: *true* if $c$ is a terminal configuration and $\Gamma$ is dynamic controllable from $c$, *false* otherwise
  **foreach** $s \in \Sigma_{\mathcal{P}}[h]$ **do**
    **if** $\mathrm{Dom}(\psi) \neq \mathcal{T}_s$ **then**
      **return** *false* ▷ Not terminal
    **if** $\psi$ *is not feasible for* $\Gamma_s$ **then**
      **return** *false* ▷ Not DC
  **return** *true*

This PSPACE algorithm actually searches for a viable dynamic execution strategy and can be easily implemented as to return the one found. However, since our algorithm works by brute force, guessing the execution times among a finite but large set of possibilities, it does not seem suitable to be applied in practice. Nevertheless, by showing that polynomial space is sufficient to solve this problem, we open up the challenge of finding a practical algorithm, that only requires polynomial memory in the worst case.

## References


[1] R. Dechter, I. Meiri, and J. Pearl, "Temporal constraint networks," *Artificial Intelligence*, vol. 49, no. 1-3, pp. 61–95, May 1991.

[2] I. Tsamardinos, T. Vidal, and M. E. Pollack, "CTP: A new constraint-based formalism for conditional, temporal planning," *Constraints*, vol. 8, no. 4, pp. 365–388, 2003.

[3] T. Vidal, "Handling contingency in temporal constraint networks: from consistency to controllabilities," *Journal of Experimental and Theoretical Artificial Intelligence*, vol. 11, no. 1, pp. 23–45, Jan. 1999.

[4] P. Morris, N. Muscettola, and T. Vidal, "Dynamic control of plans with temporal uncertainty," in *IJCAI International Joint Conference on Artificial Intelligence*, 2001, pp. 494–499.

[5] C. Comin and R. Rizzi, "Dynamic consistency of conditional simple temporal networks via mean payoff games: a singly-exponential time DC-checking," *arXiv:1505.00828 [cs]*, May 2015.

[6] A. Cimatti, L. Hunsberger, A. Micheli, R. Posenato, and M. Roveri, "Sound and complete algorithms for checking the dynamic controllability of temporal networks with uncertainty, disjunction and observation," in *2014 21st International Symposium on Temporal Representation and Reasoning*. IEEE, Sep. 2014, pp. 27–36.

[7] L. Hunsberger, R. Posenato, and C. Combi, "A Sound-and-Complete Propagation-Based Algorithm for Checking the Dynamic Consistency of Conditional Simple Temporal Networks," in *2015 22nd International Symposium on Temporal Representation and Reasoning (TIME)*. IEEE, Sep. 2015, pp. 4–18.

[8] ——, "The dynamic controllability of conditional STNs with uncertainty," pp. 1–8, 2012.

[9] C. Combi, L. Hunsberger, and R. Posenato, "An algorithm for checking the dynamic controllability of a conditional simple temporal network with uncertainty," *Evaluation*, 2013.

[10] ——, "An algorithm for checking the dynamic controllability of a conditional simple temporal network with uncertainty - revisited," in *International Conference on Agents*, 2014, pp. 314–331.